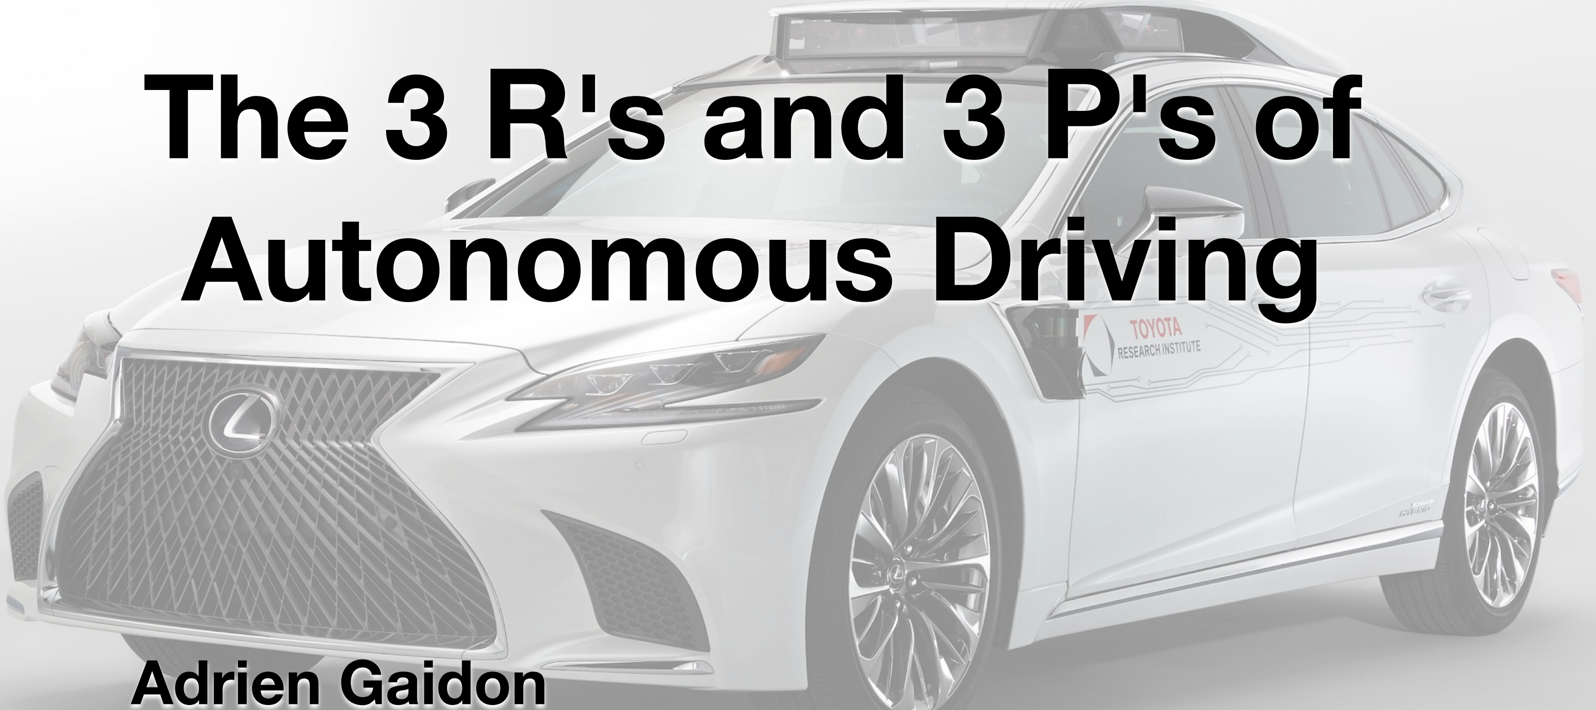

The 3 R's and 3 P's of Autonomous Driving

Adrien Gaidon

Senior Manager, Machine Learning Research
Toyota Research Institute (TRI), CA, USA

adriengaidon.com

ML-Research team (Jie Li, V. Guizilini, R. Ambrus, W. Kehl, D. Park, *et al*)
ML-Engineering team (S. Pillai, KH. Lee, A. Raventos, A. Bhargava, C. Fang, *et al*)
Wolfram Burgard, Guy Rosman,
Stanford, **TRI-AD**, **PFN**

TOYOTA RESEARCH INSTITUTE

1.35 MILLION

ROAD TRAFFIC DEATHS PER YEAR

World Health Organization: Global status report on road safety 2018

Robot = Complex Sensorimotor Loop

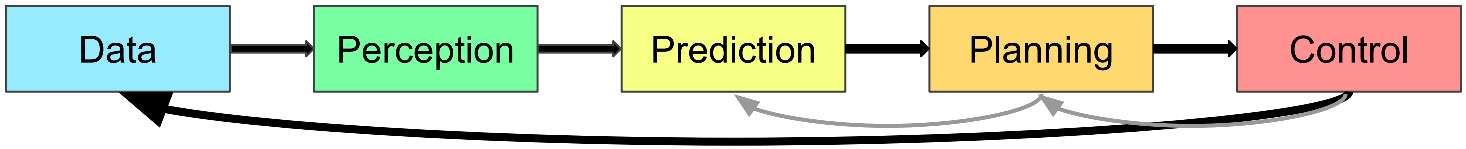

Conway's Law (paraphrased)

A system's design is isomorphic to the structure of the field.

Leaky Abstraction Law

All non-trivial abstractions, to some degree, are leaky.

Corollary: **The Arrows are Performance Bottlenecks.**

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Robustness in Perception: Data

Assumption: no expectation for a *robot* to *recognize* something *completely new*.

Domain Coverage: World-scale Fleet Learning

Problem: cannot be supervised (too much data)

Robustness in Perception: Data

Self-Supervised Learning: Geometry as Supervision

SuperDepth: Self-Supervised, Super-Resolved Monocular Depth Estimation, S. Pillai, R. Ambrus, A. Gaidon, **ICRA'19**
Robust Semi-Supervised Monocular Depth Estimation with Reprojected Distances, V. Guizilini *et al*, **CoRL'19**
Two Stream Networks for Self-Supervised Ego-Motion Estimation, R. Ambrus *et al*, **CoRL'19**
Semantically-Guided Representation Learning for Self-Supervised Monocular Depth, V. Guizilini *et al*, **ICLR'20**
3D Packing for Self-Supervised Monocular Depth Estimation**,** V. Guizilini, R. Ambrus *et al*, **CVPR'20 (oral)**
Neural Ray Surfaces for Self-Supervised Learning of Depth and Ego-motion, I. Vasiljevic, V. Guizilini *et al,* **3DV (oral)**

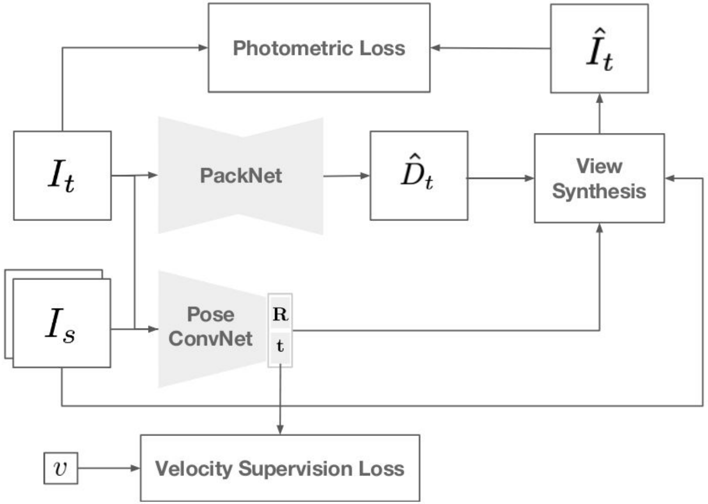

$$\hat{I}_t(p_t) = I_s(p_s) \quad p_s \sim K\hat{T}_{t\to s}\hat{D}_t(p_t)K^{-1}p_t$$

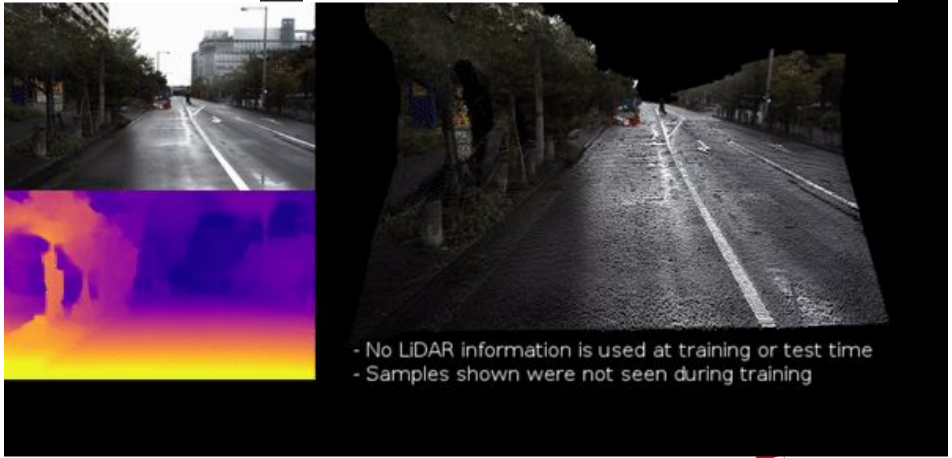

https://github.com/TRI-ML/packnet-sfm

TOYOTA RESEARCH INSTITUTE

Robustness in Perception: Data

Vision as Inverse Graphics: Auto-Labeling via Differentiable Rendering

ROI-10D: Monocular Lifting of 2D Detection to 6D Pose and Metric Shape, F. Manhardt, W. Kehl, A. Gaidon, **CVPR'19**
Autolabeling 3D Objects with Differentiable Rendering of SDF Shape Priors, S. Zakharov, W. Kehl *et al,* **CVPR'20 (oral)**
Self-Supervised Differentiable Rendering for Monocular 3D Object Detection, D. Beker, H. Kato *et al*, **ECCV'20**
Differentiable Rendering: A Survey, H. Kato, D. Beker *et al*, **arxiv:2006.12057, 2020**

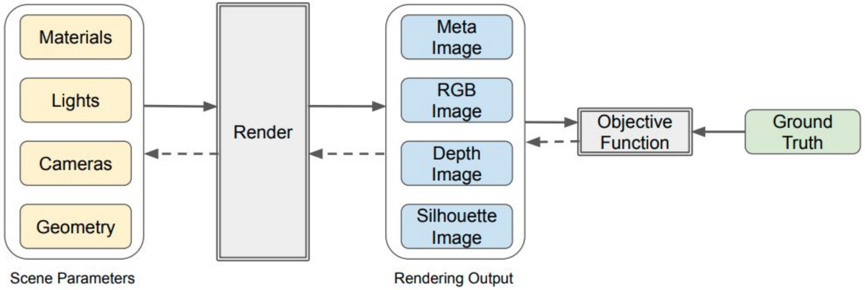

Optimization using a Differentiable Renderer

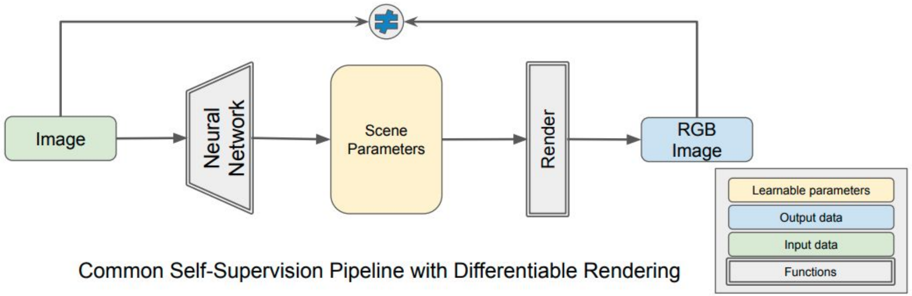

Common Self-Supervision Pipeline with Differentiable Rendering

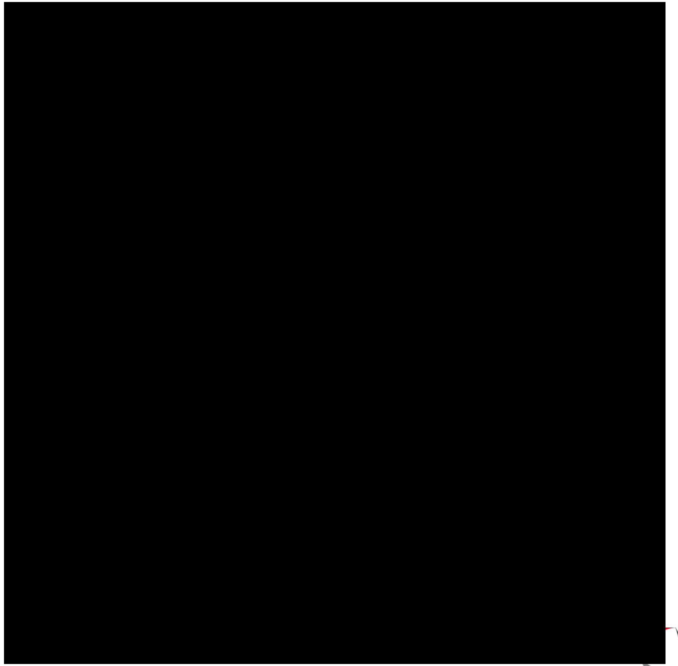

Robustness in Perception: Data

Beware of Bias: Adaptive Regularization of the Long Tail

Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss, K. Cao *et al*, **NeurIPS'19**
Heteroskedastic and Imbalanced Deep Learning with Adaptive Regularization, K. Cao *et al*, **ICML'20 w**

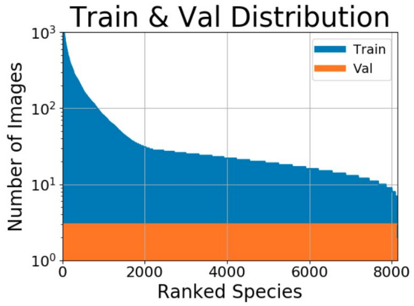

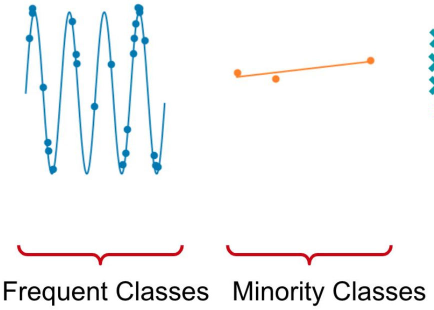

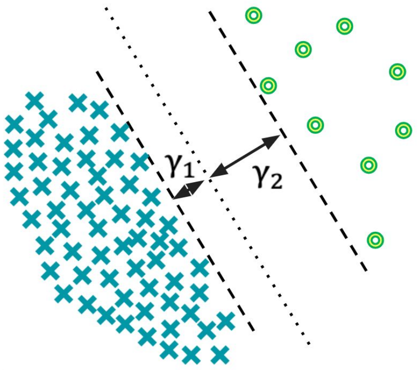

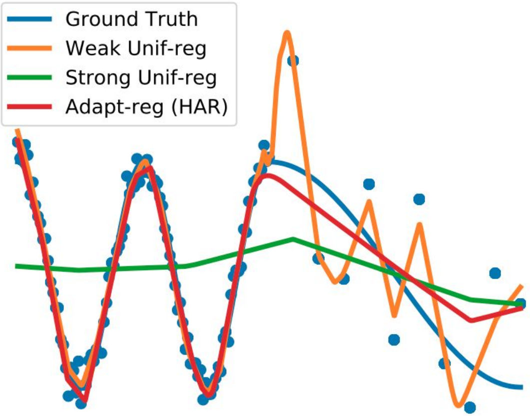

Robustness in Perception: Redundancy

Redundancy: sensors lie (Byzantine Generals)

→ Combine Radar + Lidar + Cameras + IMU + ...

Bottleneck: 3D Detection with Cameras (< Lidar)

Approach:

- Cross-sensor Auto-Labeling (2D-3D)
- Self/Semi-Supervised "pseudo-lidar"

Robust Semi-Supervised Monocular Depth Estimation with Reprojected Distances, **CoRL'19**
3D Packing for Self-Supervised Monocular Depth Estimation, **CVPR'20 (oral)**
Autolabeling 3D Objects with Differentiable Rendering of SDF Shape Priors, **CVPR'20 (oral)**

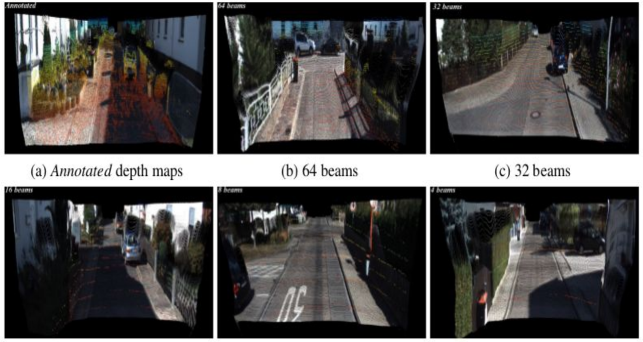

(a) *Annotated* depth maps (b) 64 beams (c) 32 beams (d) 16 beams (e) 8 beams (f) 4 beams

Robustness in Perception: Efficiency

Many sensors + complex online fusion

→ Deployment challenges (runtime, energy, stability...)

Efficiency: Core to Robustness!

Approach:

- Hardware optimization

- Sharing Computations

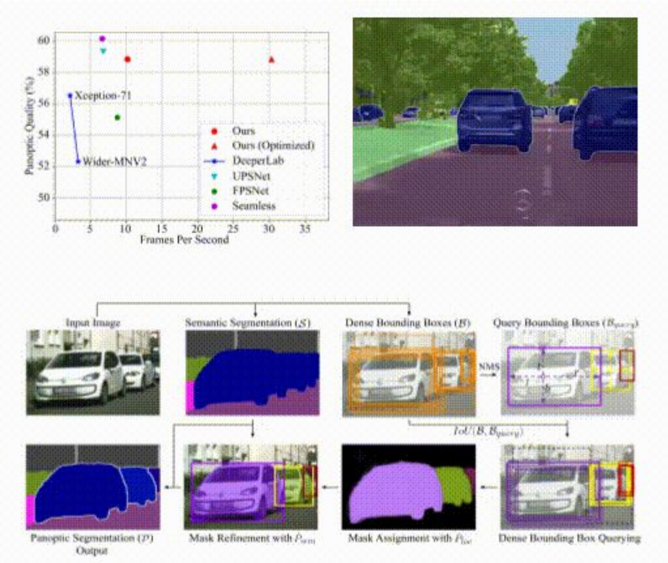

Real-Time Panoptic Segmentation from Dense Detections, R. Hou, J. Li *et al,*
CVPR'20 (oral)

3 R & 3 P of Autonomy

Robustness in Perception

Data + Redundancy + Efficiency

Randomness in Prediction

Risk-awareness in Planning

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Randomness in Prediction: Intent

Human Intent: *latent* yet governs Action

Approach: infer from Interactions and Context

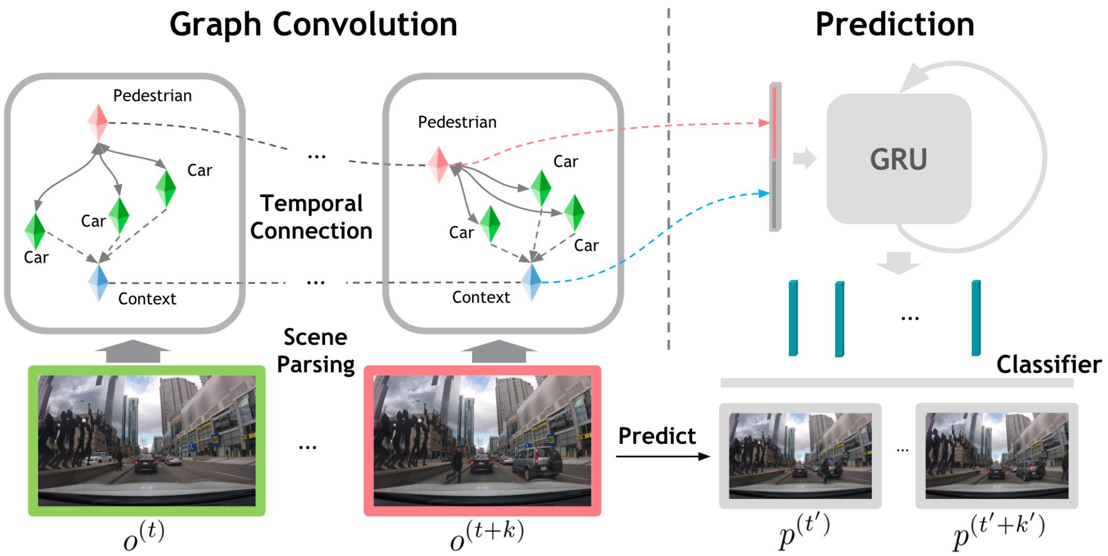

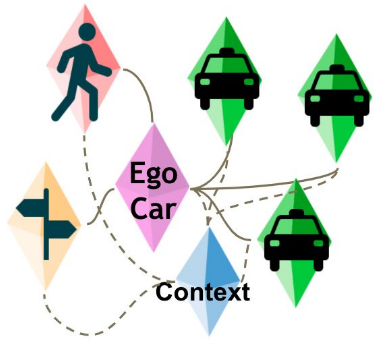

Spatiotemporal Relationship Reasoning for Pedestrian Intent Prediction, **RA-L & ICRA'20**

STIP: Stanford-TRI Intent Prediction
http://stip.stanford.edu/

Randomness in Prediction: Multi-Modality

Contingency: the Future is partly Unpredictable

Predict *Distribution* over *Plausible* Futures

"The future is already here – it's just not evenly distributed." -- W. Gibson

PREcOG: PREdiction Conditioned On Goals in Visual Multi-Agent Settings, Nicholas Rhinehart, Rowan McAllister, Kris Kitani, Sergey Levine, ICCV'19

Probabilistic Future Prediction for Video Scene Understanding, Anthony Hu, Fergal Cotter, Nikhil Mohan, Corina Gurau, Alex Kendall, arxiv:2003.06409

***Trajectron++**: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data*, Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, Marco Pavone, ECCV'20

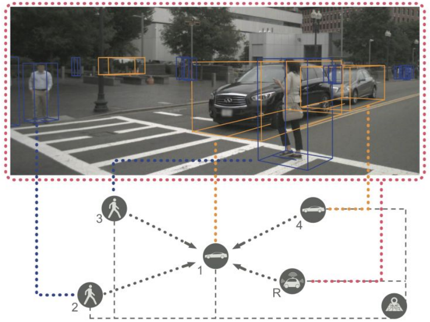

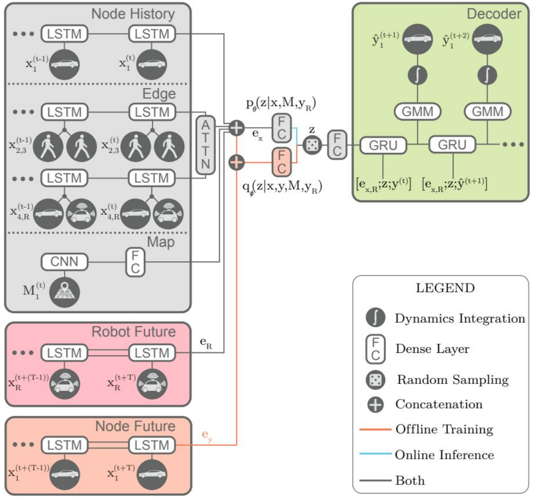

TOYOTA RESEARCH INSTITUTE

Randomness in Prediction: Multi-Modality

Multi-modality + Intent: use Human Decision-Making Process

It Is Not the Journey but the Destination: Endpoint Conditioned Trajectory Prediction,
K. Mangalam, H. Girase, S. Agarwal, K-H. Lee, E. Adeli, J. Malik, A. Gaidon, **ECCV 2020 (oral)**

Learn distribution of plausible destinations

Condition trajectory forecasting on sampled endpoints

Plan following social norms

Randomness in Prediction: Multi-Modality

Predicted-Endpoint Conditioned Network (PECNet) Architecture

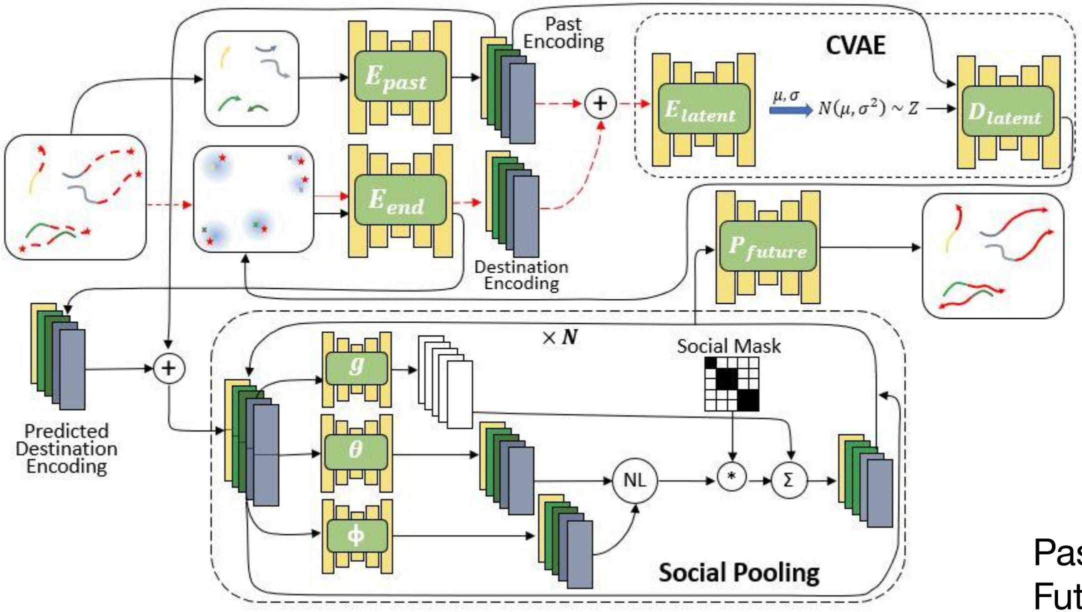

Past: 8 frames (3.2s)
Future: 12 frames (4.8s)

Randomness in Prediction: Multi-Modality

PECNet improves SOTA on ***SDD*** & UTH/ECY by **20-50%**

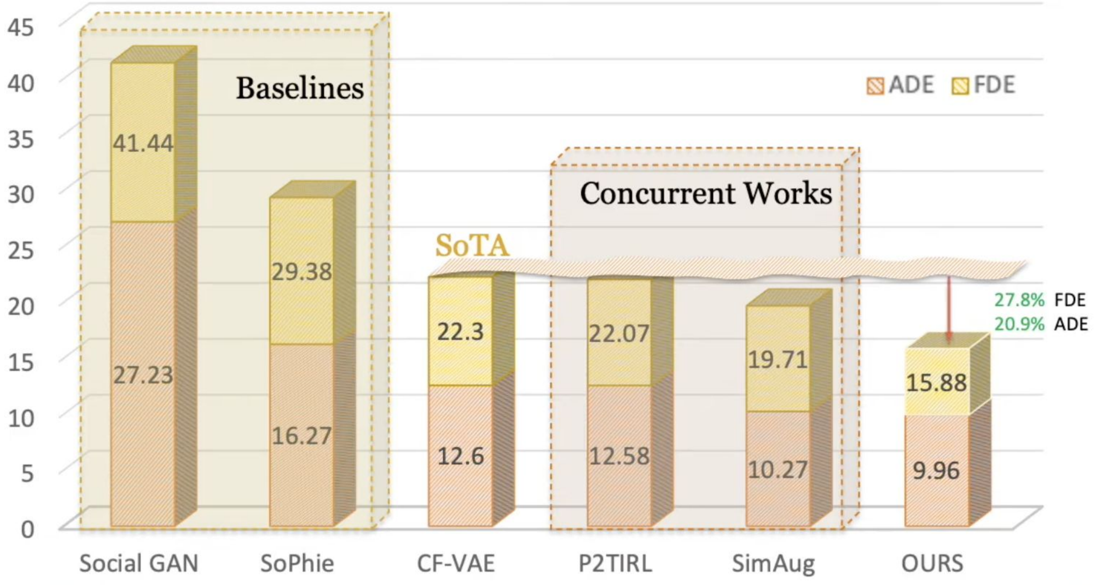

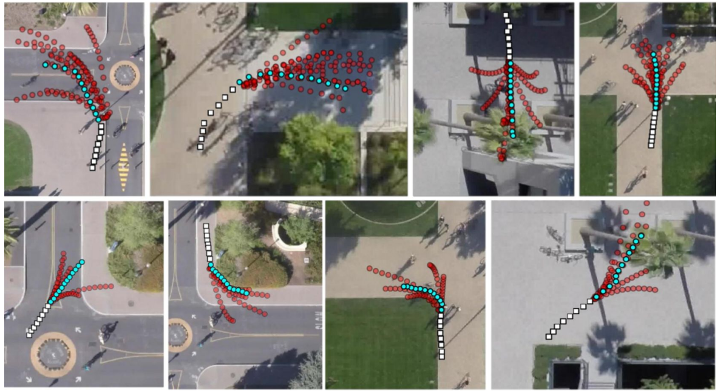

TOYOTA RESEARCH INSTITUTE

Randomness in Prediction: Uncertainty

Probabilistic Robotics: Every State is a Distribution!

Modular System: Uncertainty Propagates Throughout

***Uncertainty-Aware Driver Trajectory Prediction at Urban Intersections*, X Huang, S McGill, B C. Williams, L Fletcher, G Rosman, ICRA'19**

Probabilistic 3D Multi-Object Tracking for Autonomous Driving, H Chiu, A Prioletti, J Li, J Bohg, arxiv:2001.05673 (NuScenes winner)

A General Framework for Uncertainty Estimation in Deep Learning, A Loquercio, M Segù, D Scaramuzza, RA-L & ICRA'20

Uncertainty Estimation Using a Single Deep Deterministic Neural Network, J van Amersfoort, L Smith, YW Teh, Y Gal, arxiv:2003.02037

Can Autonomous Vehicles Identify, Recover From, and Adapt to Distribution Shifts? A Filos, P Tigas, R McAllister, N Rhinehart, S Levine, Y Gal, ICML'20

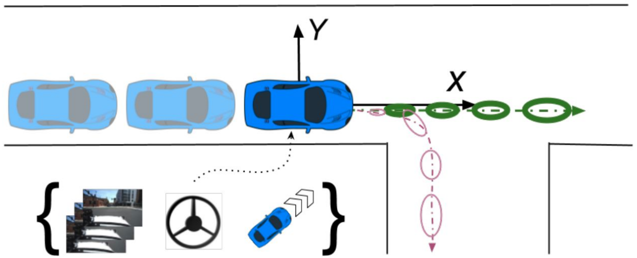

Randomness in Prediction: Uncertainty

Integrating Prediction in Planning & Control

Risk-Sensitive Sequential Action Control with Multi-Modal Human Trajectory Forecasting for Safe Crowd-Robot Interaction, H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, M. Schwager, **IROS'20**

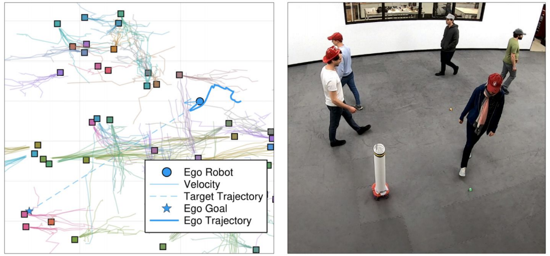

Fig. 1: The proposed RSSAC-Trajectron++ framework is effective for safe robot navigation in a social environment densely populated with humans. (Left) A simulation environment with real human trajectories from the UCY/UNIV scene [8], overlaid with predictions sampled from Trajectron++. (Right) A risk-sensitive robot safely navigating itself alongside 5 humans.

MATS: An Interpretable Trajectory Forecasting Representation for Planning and Control, B. Ivanovic, A. Elhafsi, G. Rosman, A. Gaidon, M. Pavone,

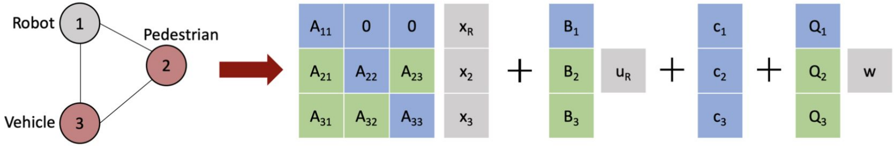

Figure 1: Our model represents edges in a spatiotemporal graph as blocks of dynamical system matrices. Blocks that are fully determined by dynamics (e.g., A's diagonal and c) or solely involve the ego-robot (e.g., the first block of B) are not learned (blue). All other blocks model agent-agent interactions or uncertainty, and are learned (green). Note that off-diagonal blocks in the ego-robot's row in A are zero, encoding the fact that the ego-robot is only controlled by its actions, $\mathbf{u}_R$.

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Intent + Multi-modality + Uncertainty

Risk-awareness in Planning

TOYOTA RESEARCH INSTITUTE

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Risk-awareness in Planning: Imitation

Demonstrations (good or bad): *trillions* of km/year!

Imitation Learning: simple, scalable, end-to-end

Exploring the Limitations of Behavior Cloning for Autonomous Driving,
F. Codevilla, E. Santana, AM. Lopez, A. Gaidon, ICCV'19 (oral)

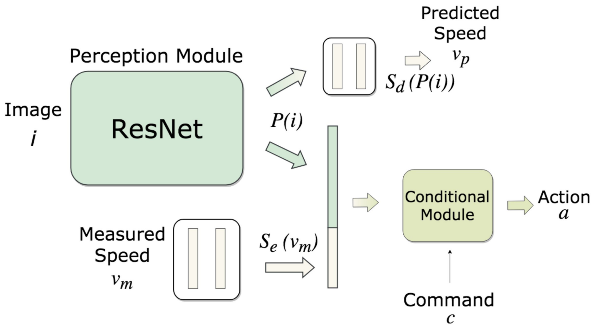

Bigger models, pretraining, more data helps... but

Dataset Biases & Variance issues

	Task	Variance
CILRS	Empty	23%
	Regular	26%
	Dense	42%
CILRS (ImageNet)	Empty	4%
	Regular	12%
	Dense	38%

Risk-awareness in Planning: Imitation

Modularity thanks to Computer Vision

(Does computer vision matter for action? B Zhou, P Krähenbühl, V Koltun, Science Robotics 2019)

But Safe Modular System → False Positives

Learn to plan under perception/prediction uncertainty

Driving Through Ghosts: Behavioral Cloning with False Positives,

A. Bühler, A. Gaidon, A. Cramariuc, R. Ambrus, G. Rosman, W. Burgard, IROS'20 (to appear)

Risk-awareness in Planning: Imitation

Modularity thanks to Computer Vision

(Does computer vision matter for action? B Zhou, P Krähenbühl, V Koltun, Science Robotics 2019)

But Safe Modular System → False Positives

Learn to plan under perception/prediction uncertainty

Driving Through Ghosts: Behavioral Cloning with False Positives,
A. Bühler, A. Gaidon, A. Cramariuc, R. Ambrus, G. Rosman, W. Burgard, IROS'20 (to appear)

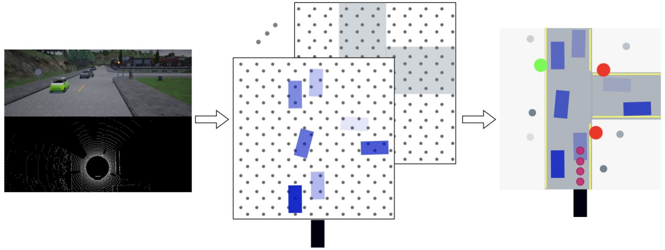

TOYOTA RESEARCH INSTITUTE

Risk-awareness in Planning: Near-Accidents

Imitation in Near-Accidents? Phase Transitions

Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving, Z. Cao, E. Biyik, W. Z. Wang, A. Raventos, A. Gaidon, G. Rosman, D. Sadigh, RSS'20

RL to switch between basic IL policies

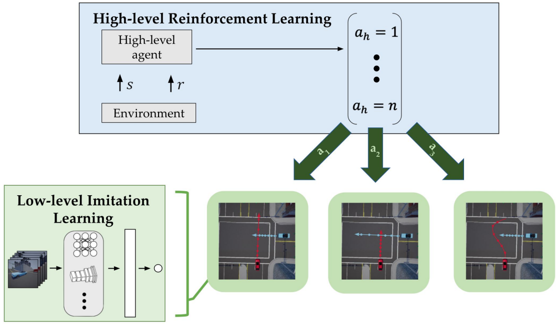

TOYOTA RESEARCH INSTITUTE

Risk-awareness in Planning: Near-Accidents

Imitation in Near-Accidents? Phase Transitions

Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving,
Z. Cao, E. Biyik, W. Z. Wang, A. Raventos, A. Gaidon, G. Rosman, D. Sadigh, RSS'20

RL to switch between basic IL policies

Less collisions + human-like

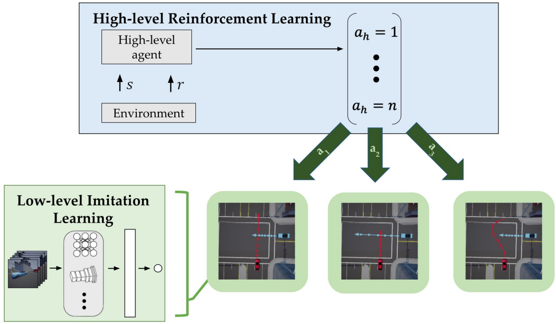

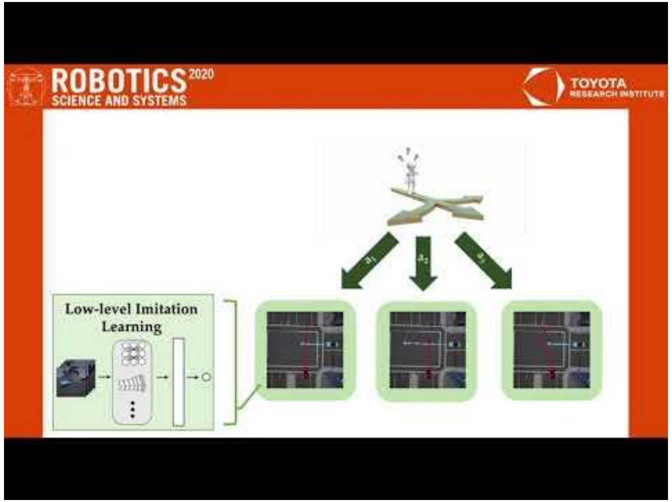

Risk-awareness in Planning: Safety

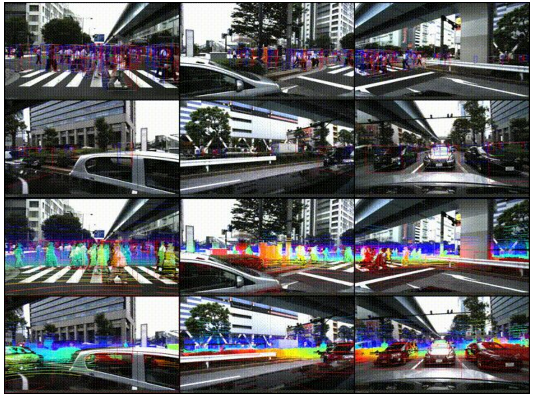
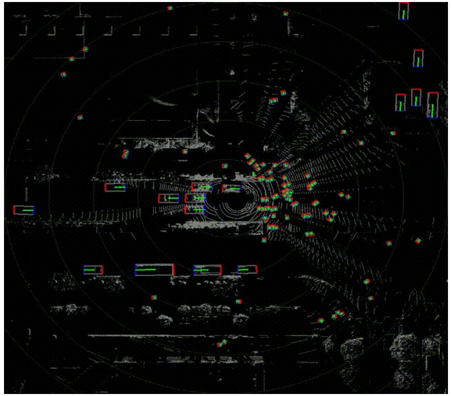

Perception / Prediction: always stochastic

Basis for *Safety-critical* Decisions in *Real-time*?

Safe Autonomy requires *Risk-Awareness*

Risk-awareness in Planning: Safety

Risk-Sensitive* Sequential Action Control with *Multi-Modal Human Trajectory Forecasting* for *Safe Crowd-Robot Interaction, H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, M. Schwager, **IROS'20**

Environment

Probabilistic Human Motion Prediction

Model-based Planning / Control

Control

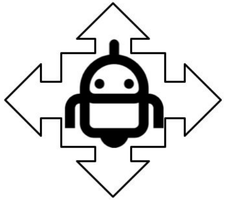

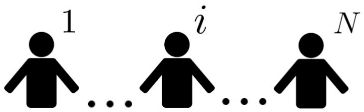

$$p^i_{k+1} = p^i_k + y^i_k \quad \{y^{1:N}_{1:T}\} \sim \mathcal{D}$$

- **Discrete-Time**
- **Stochastic**
- Arbitrary Distribution

Entropic Risk

$$R_{\mathcal{D},\sigma}(J) \triangleq \frac{1}{\sigma} \log \left(\mathbb{E}_{\mathcal{D}}[e^{\sigma J}] \right)$$

$$R_{\mathcal{D},\sigma}(J) \approx \mathbb{E}_{\mathcal{D}}[J] + \frac{\sigma}{2} \mathrm{Var}_{\mathcal{D}}(J)$$

Risk-neutral

Risk-sensitive

$\sigma = 0$

$\sigma > 0$

$$\dot{x}(t) = f\left(x(t)\right) + H\left(x(t)\right) u(t)$$

- **Continuous-Time**
- Deterministic
- Control-Affine

Risk-awareness in Planning: Safety

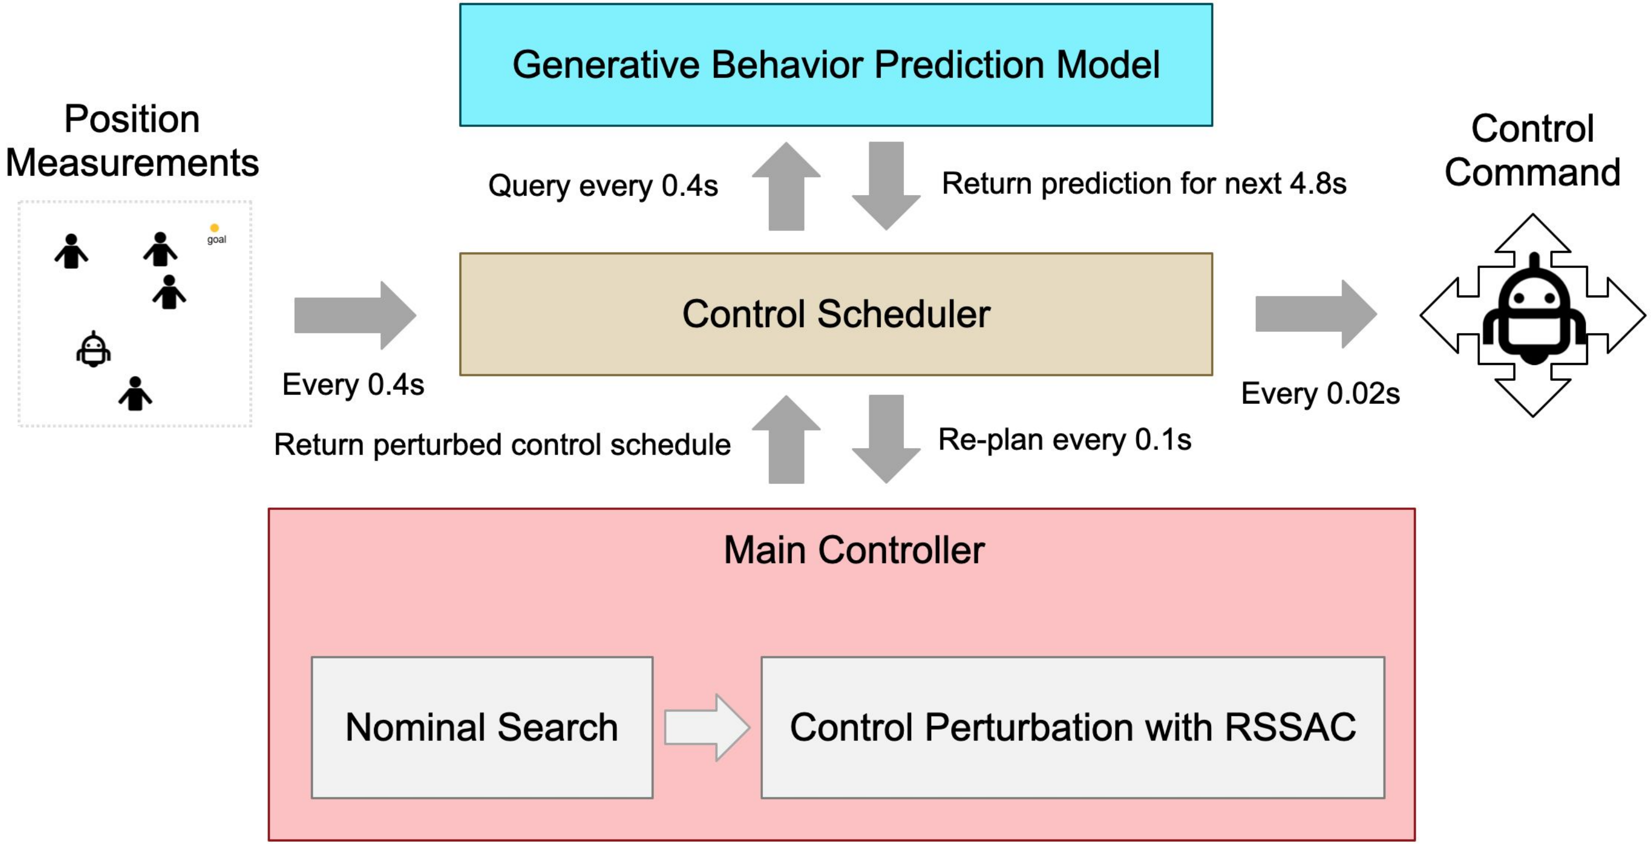

Risk-awareness in Planning: Safety

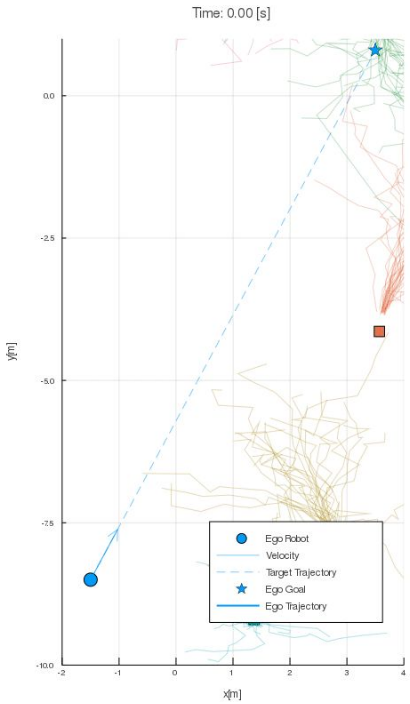

σ = 0.0 (Risk-Neutral)

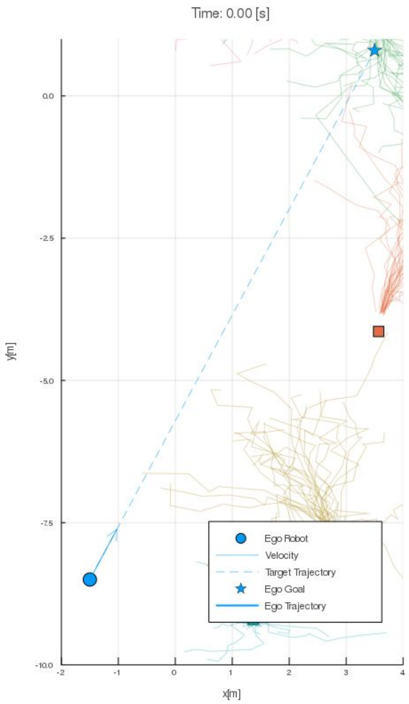

σ = 1.0 (Risk-Sensitive)

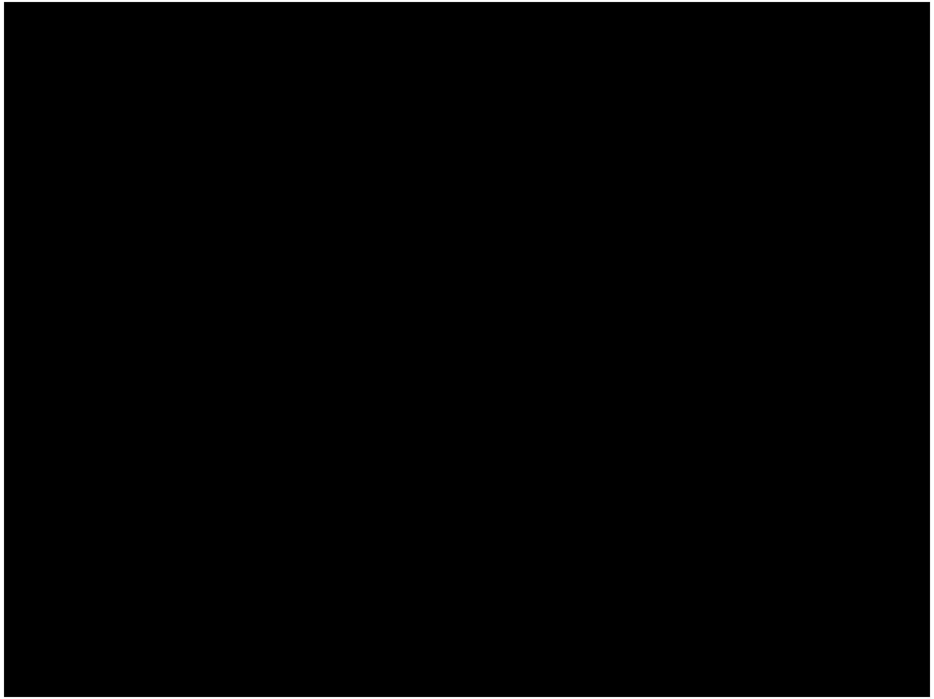

Game-Theoretic Planning for Risk-Aware Interactive Agents, M. Wang, N. Mehr, A. Gaidon, M. Schwager **IROS'20**

TOYOTA RESEARCH INSTITUTE

Risk-awareness in Planning: Causality

Decision-making: beyond pure data → Causal Inference

Counterfactuals in sim: find planner bugs and fixes

Behaviorally Diverse Traffic Simulation via Reinforcement Learning, S. Maruyama *et al*, **IROS'20**

Discovering Avoidable Planner Failures of Autonomous Vehicles using Counterfactual Analysis in Behaviorally Diverse Simulation, D. Nishiyama *et al*, **ITSC'20**

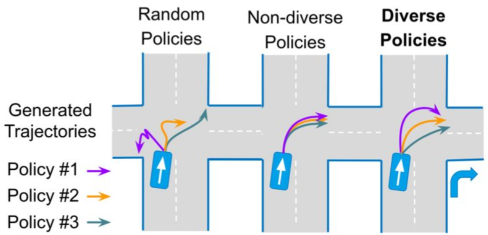

Diversity	High	Low	**High**
Driving Skills	Low	High	**High**

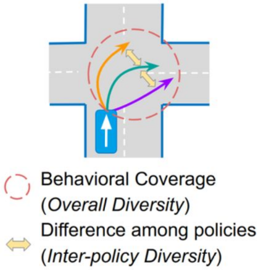

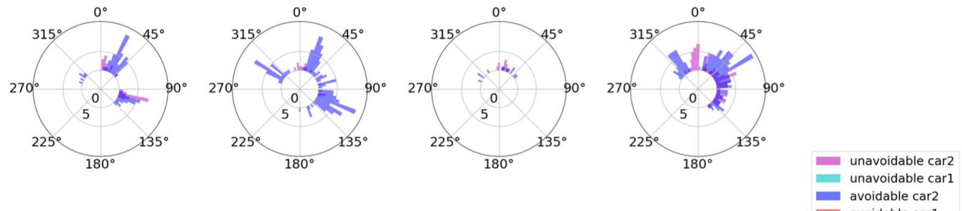

(a) LessDiverse, Right-turn (b) Diverse, Right-turn (c) LessDiverse, Crossing (d) Diverse, Crossing

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Imitation + Safety + Causality

Robot = Complex Sensorimotor Loop

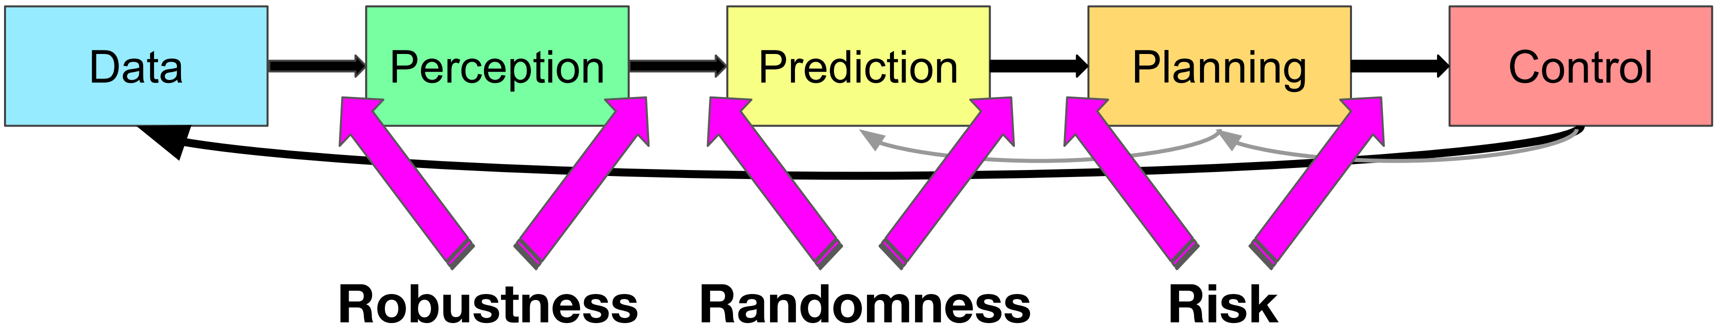

Law of Arrows: In a modular sensorimotor system, the performance bottlenecks are at the *interface* between modules.

Work on the Arrows!

Robust Perception: Data + Redundancy + Efficiency

Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss, **NeurIPS'19**
SuperDepth: Self-Supervised, Super-Resolved Monocular Depth Estimation, **ICRA'19**
Robust Semi-Supervised Monocular Depth Estimation with Reprojected Distances, **CoRL'19**
Two Stream Networks for Self-Supervised Ego-Motion Estimation, **CoRL'19**
Semantically-Guided Representation Learning for Self-Supervised Monocular Depth, **ICLR'20**
3D Packing for Self-Supervised Monocular Depth Estimation, **CVPR'20 (oral)**
Real-Time Panoptic Segmentation from Dense Detections, **CVPR'20 (oral)**
Autolabeling 3D Objects with Differentiable Rendering of SDF Shape Priors, **CVPR'20 (oral)**
Self-Supervised Differentiable Rendering for Monocular 3D Object Detection, **ECCV'20**
PillarFlow: End-to-end Birds-eye-view Flow Estimation for Autonomous Driving, **IROS'20**
Neural Ray Surfaces for Self-Supervised Learning of Depth and Ego-motion, **3DV (oral)**

Random Prediction: Intent + Multi-modality + Uncertainty

Spatiotemporal Relationship Reasoning for Pedestrian Intent Prediction, **RA-L & ICRA'20**
It Is Not the Journey but the Destination: Endpoint Conditioned Trajectory Prediction, **ECCV'20 (oral)**
MATS: An Interpretable Trajectory Forecasting Representation for Planning and Control, arxiv:2009.07517

Risk-aware Planning: Imitation + Safety + Causality

Exploring the Limitations of Behavior Cloning for Autonomous Driving, **ICCV'19 (oral)**
Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving, **RSS'20**
Driving Through Ghosts: Behavioral Cloning with False Positives, **IROS'20**
Risk-Sensitive Sequential Action Control with Multi-Modal Human Trajectory Forecasting [...], **IROS'20**
Game-Theoretic Planning for Risk-Aware Interactive Agents, **IROS'20**
Behaviorally Diverse Traffic Simulation via Reinforcement Learning, **IROS'20**
Discovering Avoidable Planner Failures [...] in Behaviorally Diverse Simulation, **ITSC'20**

TRI Dataset Releases

STIP: Stanford-TRI Intent Prediction
http://stip.stanford.edu/
DDAD: Dense Depth for AD
https://github.com/TRI-ML/DDAD

Workshops co-organized by TRI

ICML: AI for AD (AIAD)
https://sites.google.com/view/aiad2020
ECCV: Perception for AD (PAD)
https://sites.google.com/view/pad2020

Thanks to many collaborators!
E. Adeli, S. Agarwal, R. Ambrus, T. Ando, N. Arechiga, D. Beker, A. Bhargava, E. Biyik, W. Burgard, A. Bühler, H. Cai, K. Cao, Z. Cao, MY. Castro, F. Codevilla, A. Cramariuc, J. DeCastro, B. Douillard, C. Fang, H. Girase, V. Guizilini, K. Hamzaoui, R. Hou, DA. Huang, B. Ivanovic, H. Kato, W. Kehl, M. Kliemann, KH. Lee, J. Li, B. Liu, AM. Lopez, J. Lynch, T. Ma, J. Malik, K. Mangalam, F. Manhardt, S. Maruyama, T. Matsuoka, N. Mehr, MA. Morariu, JC Niebles, H. Nishimura, D. Nishiyama, Y. Ouyang, B. Pan, JM. Pan, M. Pavone, S. Pillai, A. Raventos, G. Ros, G. Rosman, D. Sadigh, E. Santana, M. Schwager, A. Shenoi, S. Shiroshita, T. Tagawa, M. Wang, WZ. Wang, C. Wei, S. Zakharov

TOYOTA RESEARCH INSTITUTE